

PRECONDITIONING THE PRIOR TO OVERCOME SATURATION IN BAYESIAN INVERSE PROBLEMS

SERGIOS AGAPIOU AND PETER MATHE

ABSTRACT. We study Bayesian inference in statistical linear inverse problems with Gaussian noise and priors in Hilbert space. We focus our interest on the posterior contraction rate in the small noise limit. Existing results suffer from a certain saturation phenomenon, when the data generating element is too smooth compared to the smoothness inherent in the prior. We show how to overcome this saturation in an empirical Bayesian framework by using a non-centered data-dependent prior. The center is obtained from a preconditioning regularization step, which provides us with additional information to be used in the Bayesian framework. We use general techniques known from regularization theory. To highlight the significance of the findings we provide several examples. In particular, our approach allows to obtain and, using preconditioning improve after saturation, minimax rates of contraction established in previous studies. We also establish minimax contraction rates in cases which have not been considered so far.

1. SETUP

We consider the following linear equation in real Hilbert space

$$y^\delta = Kx + \delta\eta,$$

where $K: X \rightarrow Y$ is a linear operator acting between the real separable Hilbert spaces X and Y , $\eta \sim \mathcal{N}(0, \Sigma)$ is an additive centered Gaussian noise, and $\delta > 0$ is a scaling constant modelling the size of the noise. Here, the covariance operator $\Sigma: Y \rightarrow Y$ is a self-adjoint and positive definite bounded linear operator. We formally pre-whiten this equation and get

$$z^\delta = \Sigma^{-1/2}y^\delta = \Sigma^{-1/2}Kx + \delta\xi,$$

where now $\xi \sim \mathcal{N}(0, I)$ is Gaussian white noise. We assign $T := \Sigma^{-1/2}K$, and we assume that this is bounded by imposing the condition $\mathcal{R}(K) \subset \mathcal{D}(\Sigma^{-1/2})$. We hence arrive to the data model

$$(1) \quad z^\delta = Tx + \delta\xi,$$

Date: Wednesday 24th September, 2014: 01:00.

2010 *Mathematics Subject Classification.* 62G20, secondary: 62C10, 62F15, 45Q05.

Key words and phrases. Bayesian inverse problem, posterior contraction, saturation.

and we consider the Bayesian approach to the statistical inverse problem of finding x from the observation z^δ . We assume Gaussian priors on x , distributed according to $\mathcal{N}(0, \frac{\delta^2}{\alpha} C_0)$, where $C_0 : X \rightarrow X$ is a positive definite, self-adjoint and trace class linear operator, and $\alpha > 0$ is a scaling constant. Linearity suggests that the posterior is also Gaussian and in this paper we are interested in the asymptotic performance of the posterior in the small noise limit, $\delta \rightarrow 0$.

Squared posterior contraction. Suppose that we observe data z^δ generated from the model (1) for a fixed underlying true element $x^* \in X$ and corresponding to a noise level δ . It is then reasonable to expect that for small δ and for appropriate values of α , the posterior Gaussian distribution will concentrate around the true data-generating element x^* . As we discuss below, this concentration will be driven by the *squared posterior contraction* (SPC), given as

$$(2) \quad \text{SPC} := \mathbb{E}^{x^*} \mathbb{E}_\alpha^{z^\delta} \|x^* - x\|^2,$$

where the outward expectation is taken with respect to the data generating distribution, that is, the distribution generating z^δ when x^* is given, and the inward expectation is taken with respect to the posterior distribution, given data z^δ and having chosen a parameter α . The Gaussian posterior distribution has a posterior mean, say $x_\alpha^\delta = x_\alpha^\delta(z^\delta; \alpha)$, and a posterior covariance, say $C^\delta(\alpha)$, which is independent from the data z^δ , and thus deterministic. Then the inner expectation obeys the usual bias-variance decomposition

$$\mathbb{E}_\alpha^{z^\delta} \|x^* - x\|^2 = \|x^* - x_\alpha^\delta\|^2 + \text{tr} [C^\delta(\alpha)].$$

Applying the expectation with respect to the data generating distribution, we obtain that

$$\mathbb{E}^{x^*} \mathbb{E}_\alpha^{z^\delta} \|x^* - x\|^2 = \mathbb{E}^{x^*} \|x^* - x_\alpha^\delta\|^2 + \text{tr} [C^\delta(\alpha)].$$

The quantity $\mathbb{E}^{x^*} \|x^* - x_\alpha^\delta\|^2$ represents the mean integrated squared error (MISE) of the posterior mean viewed as an estimator of x^* , and it has again a bias-variance decomposition into squared bias $b_{x^*}^2(\alpha) := \|x^* - \mathbb{E}^{x^*} x_\alpha^\delta\|^2$ and estimation variance $V^\delta(\alpha) := \mathbb{E}^{x^*} \|x_\alpha^\delta - \mathbb{E}^{x^*} x_\alpha^\delta\|^2$. We have thus decomposed the squared posterior contraction into respectively the *squared bias*, the *estimation variance*, and the *spread in the posterior distribution*

$$(3) \quad \text{SPC}(\alpha, \delta) = b_{x^*}^2(\alpha) + V^\delta(\alpha) + \text{tr} [C^\delta(\alpha)].$$

We emphasize here, that the decomposition remains valid in the more general case of non-centered Gaussian priors.

It is clear, that if possible the hyper-parameter α should be chosen in a way that optimizes the SPC. This raises several questions and challenges.

First, how do the estimation variance $V^\delta(\alpha)$ and the posterior spread $\text{tr}[C^\delta(\alpha)]$ relate? In previous studies, these quantities appear to be either of the same order, see proof of [8, Thm 4.1], or the posterior spread dominates the estimation variance, see proofs of [3, Thm 4.3] and [9, Thm 2.1]. As was first highlighted in [11], there is a natural relation $V^\delta(\alpha) \leq \text{tr}[C^\delta(\alpha)]$, whenever the prior is centered.

The *posterior contraction rate* is concerned with the concentration rate of the posterior distribution around the truth, in the small noise limit $\delta \rightarrow 0$, and given a prior distribution. It is well known, that the square root of the convergence rate of SPC is a posterior contraction rate (see for example [2, Section 7]). Given the prior scaling assumed here, SPC decays to zero provided that the parameter α is chosen such that $\alpha = \alpha(\delta) \rightarrow 0$ in an appropriate manner. The study of this decay was the subject of the papers [8, 2, 9, 3]. The obtained rates of convergence depend on the relationship between the regularity of the data-generating element x^* and the regularity inherent in the prior (see [5, § 2.4] for details on the regularity of draws from Gaussian measures in Hilbert space). The general message is that if the prior regularity matches the regularity of x^* , then the convergence rate of SPC is the minimax-optimal rate even without rescaling the prior, that is for the scaling considered here, α should be chosen to be equal to δ^2 . If there is a mismatch between the prior regularity and the regularity of the truth, then the minimax rate can be achieved by appropriately rescaling the prior. If the prior is smoother than the truth, then there exists an a priori parameter choice rule $\alpha = \alpha(\delta)$ such that $\frac{\delta^2}{\alpha} \rightarrow \infty$ as $\delta \rightarrow 0$, which gives the optimal rate. If however the prior is rougher than the truth, then the minimax rate can be achieved by appropriate choices $\alpha = \alpha(\delta)$ such that $\frac{\delta^2}{\alpha} \rightarrow 0$ as $\delta \rightarrow 0$, in general only up to a maximal smoothness of x^* . As quoted in [8], rescaling can make the prior arbitrarily 'rougher' but not arbitrarily 'smoother'. A closer look at the situation reveals, and we shall highlight this in our subsequent analysis, that the estimation bias, which is part of the SPC in (3), is responsible for this phenomenon. Bounds for the bias depend on the inter-relation between the underlying solution smoothness and the capability of the chosen (Tikhonov-type since we have Gaussian priors) reconstruction by means of x_α^δ to take it into account. The capability of such a scheme to take smoothness into account is called *qualification* of the scheme, whereas the limited decay rate of the bias, as $\alpha \rightarrow 0$, due to the chosen reconstruction scheme, is called *saturation* of the scheme. Details will be given below.

Finally, the optimal choices $\alpha = \alpha(\delta)$ depend on the regularity of x^* , which is in practice unknown. In the literature there have been two strategies to overcome these difficulties, both in the simplified setting of the white noise model (that is, the case $K = \Sigma = I$). The first one is to attempt to learn the correct scaling from the data, either by

using a maximum likelihood empirical Bayes approach, or by a fully hierarchical approach. This has been studied in [15], where the results show that in both approaches the minimax rate is achieved but again up to a maximal regularity of the truth (which surprisingly is smaller than the one for the oracle type choice of α). The second strategy is to not rescale the prior but rather attempt to learn the correct regularity from the data, again either using a maximum likelihood empirical Bayes or a fully hierarchical approach. This is the topic of [7], where indeed the authors show that the minimax rate is achieved by both of the approaches. The last method seems to address both the issue of saturation and of choosing α , however, all of the methods mentioned in this paragraph can be difficult to implement. On the one hand as it is shown in [1], the implementation of the hierarchical approach in non-trivial problems is problematic in high dimensions and for small noise, while on the other hand the above empirical Bayes approaches involve solving an optimization problem which also becomes difficult for non-trivial problems.

Paradigm. Here we consider the following alternative paradigm. Suppose we want to use a Gaussian prior with covariance C_0 , and prior mean m_0 to gain posterior inference for the problem (1). The question we address is whether the prior center m_0 has a significant impact on the posterior contraction rate, and if so, how to choose it 'optimally' in the presence of data. The subsequent analysis will show that the convergence rate of SPC will improve by an appropriate adjustment of the prior if the underlying solution x^* has large smoothness. In terms of the previous discussion, for a prior of fixed smoothness this enables us to make *a priori* choices of $\alpha = \alpha(\delta)$ such that the posterior contraction rate is minimax-optimal even for higher smoothness of x^* , by choosing an appropriate center m_0 of the prior distribution. The proposed re-centering $m_0 = m_0(z^\delta; \alpha)$ of the prior depends on the data z^δ and the parameter α , it is not static. However, it can easily be managed by a regularization step preprocessing the Bayes step. We anticipate these results in the following Figure 1. This figure highlights the results as described in § 4.2.

We capture the advantages in a few lines:

- the user may choose a (centered) Gaussian prior of arbitrary smoothness;
- after observing data z^δ , a prior center, say $m_0 = m_0(z^\delta; \alpha)$ is determined by some deterministic regularization;
- if this preprocessing regularization has enough qualification, then the posterior distribution will contract order optimally regardless of the solution smoothness. If not, then the contraction rate is at least as good as the rate corresponding to a centered prior.




FIGURE 1. Exponents of convergence rates of SPC plotted against Sobolev-like smoothness of the truth β , for different methods of choosing the prior mean m_α^δ , in the moderately ill-posed problem discussed in § 4.2. We set $D := 1 + 2a + 2p$, the saturation point when no preconditioning of the prior mean is used. Rates calculated for $a = 0.5$, $p = 1$.

- this preprocessing step has no effect on the parameter choice; so *any* choice $\alpha = \alpha(\delta; z^\delta)$ which yields ‘optimal’ contraction without preprocessing will retain this property, and will eventually extend this optimality property for higher solution smoothness.

Outline. In order to explain the new paradigm we first study the impact of using a non-centered prior to the posterior mean and covariance. Then we specify the prior centering by means of using a linear regularization in Eq. (6), as such is known from regularization theory. Next, we provide explicit representations of the quantities involved in the subsequent analysis, the posterior mean, the posterior covariance, and formulas for the bias and estimation variance, see Eq. (7)–(10).

The main results are given in Section 3, after confining ourselves to the case of commuting operators C_0 and T^*T , expressed in terms of a specific link condition. We first derive bounds for the estimation bias in Proposition 3.1, and these bounds are crucial for overcoming the saturation. Then we introduce the net posterior spread in § 3.3, which is the unscaled version of the posterior spread, and we highlight its properties. We then combine to obtain our main result on the convergence of SPC, which is Theorem 1.

To emphasize the significance of our results we discuss in Section 4 specific examples some of which were previously studied in [8, 3, 9]. In order to facilitate the reading of the study we postpone all proofs to the final Section 5.

2. SETTING THE PACE

As mentioned above, we shall discuss a preprocessing of the prior by choosing it non-central, that is, we will introduce a shift m_0 , such that the prior will be Gaussian with $\mathcal{N}(m_0, \frac{\delta^2}{\alpha} C_0)$. In particular, we are interested in understanding the impact of the shift m_0 on the convergence rate of SPC. For the reader's convenience, we start with deriving formulas for the posterior mean x_α^δ in this context.

We first recall the representation of the posterior mean m and posterior covariance C when a centered prior $\mathcal{N}(0, \frac{\delta^2}{\alpha} C_0)$ is used. In this case we know, see for example [12, 10], that almost surely with respect to the joint distribution of (x, z^δ) the posterior is Gaussian, $\mathcal{N}(m, C)$, for

$$(4) \quad m = C_0^{1/2} (\alpha I + B^* B)^{-1} B^* z^\delta,$$

and

$$(5) \quad C = \delta^2 C_0^{1/2} (\alpha I + B^* B)^{-1} C_0^{1/2},$$

where we define the compact operator $B := TC_0^{\frac{1}{2}}$. Re-centering the prior towards m_0 does not affect the posterior covariance C . To obtain the shift in the posterior mean we rewrite (1) as

$$z^\delta - Tm_0 = T(x - m_0) + \delta\xi$$

Thus if $x \sim \mathcal{N}(m_0, C_0)$ then $x - m_0 \sim \mathcal{N}(0, C_0)$. We are in the usual context with centered prior but new data $z^\delta - Tm_0$. This gives the representation for the posterior mean (shifting back towards m_0) as

$$\begin{aligned} x_\alpha^\delta &= m_0 + C_0^{1/2} (\alpha I + B^* B)^{-1} B^* (z^\delta - Tm_0) \\ &= C_0^{1/2} (\alpha I + B^* B)^{-1} B^* z^\delta + m_0 - C_0^{1/2} (\alpha I + B^* B)^{-1} B^* Tm_0 \\ &= C_0^{1/2} (\alpha I + B^* B)^{-1} B^* z^\delta + C_0^{1/2} (I - (\alpha I + B^* B)^{-1} B^* B) C_0^{-1/2} m_0 \\ &= C_0^{1/2} (\alpha I + B^* B)^{-1} B^* z^\delta + C_0^{1/2} s_\alpha(B^* B) C_0^{-1/2} m_0, \end{aligned}$$

where we introduce the function $s_\alpha(t) = \alpha/(\alpha + t)$, $\alpha, t > 0$, applied to the self-adjoint operator $B^* B$ by using spectral calculus.

It is well-understood from previous Bayesian analysis that a static choice of m_0 will not have impact on the posterior contraction. However, within our new paradigm we choose any regularization scheme g_α and assign the prior center as

$$(6) \quad m_0(z^\delta; \alpha) := m_\alpha^\delta = C_0^{1/2} g_\alpha(B^* B) B^* z^\delta.$$

We introduce linear regularization schemes as follows.

Definition 1 (linear regularization). *Let $b = \|B^* B\|$. A family of piece-wise continuous functions $g_\alpha: (0, b] \rightarrow \mathbb{R}$, $\alpha > 0$, is called regularization filter with residual function $r_\alpha(t) = 1 - tg_\alpha(t)$, $\alpha, 0 < t \leq b$, if*

- (1) $\sup_{0 < t \leq b} |r_\alpha(t)| \leq \gamma_0$, for all $\alpha > 0$,
- (2) $\lim_{\alpha \rightarrow 0} r_\alpha(t) = 0$ for each $0 < t \leq b$, and
- (3) $\sup_{0 < t \leq b} |g_\alpha(t)| \leq \gamma_*/\alpha$, for all $\alpha > 0$.

The above requirements are the ones which are typically imposed on a linear regularization scheme, see for example [6].

Remark 2.1. The element $m_0(z^\delta; \alpha)$ belongs to the Cameron-Martin space of the prior, that is, the subspace $\mathcal{D}(C_0^{-\frac{1}{2}})$ of X , almost surely with respect to the joint distribution of (x, z^δ) . To see this combine the first assertion of Definition 1 with the fact that the operator $C_0^{1/2}$ is Hilbert–Schmidt. As a side remark, we mention that this means that the Gaussian prior measures corresponding to any parameter α , or even any regularization filter g_α , are absolutely continuous with respect to each other.

Remark 2.2. The last assertion in Definition 1 is actually stronger than the one required in [6], but it is a convenient strengthening, and most known regularization schemes obey this stronger bound.

Remark 2.3. We use the following convention: if no preconditioning is used, that is, if $g_\alpha(t) \equiv 0$, then we assign the constant function $r_\alpha(t) \equiv 1$, in order to simplify the comparison of the different settings. Specifically, without preprocessing we would naturally (and statically) use $m_0 := 0$ as the prior mean.

Example 1 (Tikhonov regularization). One of the commonly used regularization schemes is Tikhonov regularization, in which case the filter g_α is given as $g_\alpha(t) = 1/(\alpha + t)$, $\alpha, t > 0$. Notice that in the case $m_0 = 0$, the posterior mean as given in Eq. (4), has the form of the right hand side in Eq. (6) with g_α being the Tikhonov filter.

Remark 2.4. We fix once and for all, as above the function $s_\alpha(t) = \alpha/(\alpha + t)$, that is, the residual function for Tikhonov regularization. This is done in order to distinguish the (Tikhonov) regularization in the posterior mean due to the use of a Gaussian prior, from the chosen regularization for the prior preconditioning.

Example 2 (k -fold Tikhonov regularization). We may iterate Tikhonov regularization, starting from the trivial element $x_{0,\alpha} = 0$ as

$$x_{j,\alpha}^\delta := x_{j-1,\alpha}^\delta + (\alpha I + B^* B)^{-1} B^* (z^\delta - B x_{j-1,\alpha}^\delta), \quad j = 1, \dots, k.$$

For $k = 1$ this gives Tikhonov regularization. The resulting linear regularization is given by the function $g_{k,\alpha} := \frac{1}{t} \left(1 - \left(\frac{\alpha}{\alpha+t}\right)^k\right)$, $t > 0$, with corresponding residual function $r_{k,\alpha} = \left(\frac{\alpha}{\alpha+t}\right)^k$, $t > 0$. This regularization results in the prior center $m_\alpha^\delta = C_0^{1/2} x_{k,\alpha}^\delta = C_0^{\frac{1}{2}} g_{k,\alpha} (B^* B) B^* z^\delta$.

Example 3 (spectral cut-off, truncated SVD). This is a versatile scheme, which requires to know the singular value decomposition of the underlying operator. If this is available, then we let $g_\alpha(t) = 1/t$, for $t \geq \alpha$ and $g_\alpha(t) = 0$ else.

We summarize the previous considerations and fix the notation which will be used subsequently. Given prior mean m_α^δ from (6), we have that the posterior distribution is Gaussian with posterior mean, denoted as x_α^δ , given as

$$(7) \quad x_\alpha^\delta = C_0^{1/2} (\alpha I + B^* B)^{-1} B^* z^\delta + C_0^{1/2} s_\alpha(B^* B) C_0^{-1/2} m_\alpha^\delta,$$

and posterior covariance $C := C^\delta(\alpha)$ with

$$(8) \quad C^\delta(\alpha) = \delta^2 C_0^{1/2} (\alpha I + B^* B)^{-1} C_0^{1/2}.$$

Since we aim at controlling the squared posterior contraction, we have that the spread is given as $\text{tr}[C^\delta(\alpha)]$ and we next give expressions for the corresponding estimation bias and estimation variance.

Lemma 2.1. *Let x_α^δ be as in (7). Then the estimation bias and estimation variances, with posterior mean as estimator, are*

$$(9) \quad b_{x^*}(\alpha) = \left\| C_0^{1/2} s_\alpha(B^* B) r_\alpha(B^* B) C_0^{-1/2} x^* \right\|, \quad \alpha > 0,$$

and

$$(10) \quad V^\delta(\alpha) = \delta^2 \text{tr} \left[(I + \alpha g_\alpha(B^* B))^2 (\alpha I + B^* B)^{-2} B^* B C_0 \right], \quad \alpha > 0,$$

respectively.

Proposition 2.1. *Let the prior center be obtained from any regularization (with corresponding constant γ_*). Then we have that*

$$(11) \quad V^\delta(\alpha) \leq (1 + \gamma_*)^2 \text{tr}[C^\delta(\alpha)].$$

Consequently we have that

$$\mathbb{E}^{x^*} \|x^* - x_\alpha^\delta\|^2 \leq \text{SPC}(\alpha, \delta) \leq b_{x^*}^2(\alpha) + (1 + (1 + \gamma_*)^2) \text{tr}[C^\delta(\alpha)].$$

Remark 2.5. The above analysis extends the previous bound from [11, Eq. (12)] to the present context (note that without preprocessing we have that $\gamma_* = 0$). We also note that the decay of the squared posterior contraction cannot be faster than the minimax error for statistical estimation.

We thus have that in order to (asymptotically) bound the squared posterior contraction, we only need to establish bounds for the bias and the posterior spread.

3. ASSUMPTIONS AND MAIN RESULTS

We are now ready to present our main results. Before we do so, in § 3.1 we introduce several concepts used in our formulation. First, we introduce *link conditions*, relating the two operators appearing in the setting at hand. Then we introduce *source sets*, which we use for expressing the regularity of the truth. Finally, we introduce the *qualification* of a regularization which quantifies its capability to take high smoothness into account. We then present our bounds for the bias, the posterior spread and finally the squared posterior contraction in § 3.2, § 3.3 and § 3.4, respectively.

3.1. Link conditions, source sets and qualification. We call a function $\varphi: (0, \infty) \rightarrow \mathbb{R}^+$ an *index function* if it is a continuous non-decreasing function which can be extended to take the value zero at the origin.

Remark 3.1. The property of interest of an index function is its asymptotic behaviour near the origin. In some cases the 'native' index function is not defined on $(0, \infty)$, but only on some sub-interval, say $(0, \bar{t})$. Consider for example the logarithmic function $\varphi(t) = \log^{-\mu}(1/t)$, $0 < t < \bar{t} = 1$ with $\phi(0) = 0$. Then one can extend the function ϕ at some interior point $0 < t_0 < \bar{t}$ in an increasing way, for instance as $\varphi(t) = \varphi(t_0) + (t - t_0)$, $t \geq t_0$. By doing so we ensure that the extended function shares the same asymptotic properties near zero, that is, as $t \searrow 0$. In all subsequent (asymptotic) considerations it suffices to have such extensions, and this will not be mentioned explicitly.

To simplify the outline of the study we confine ourselves to commuting operators C_0 and T^*T . Specifically we do this as follows.

Assumption 3.1 (link condition). There is an index function ψ such that

$$(12) \quad \psi^2(C_0) = T^*T.$$

Along with the function ψ we introduce the function

$$(13) \quad \Theta_\psi(t) := \sqrt{t}\psi(t), \quad t > 0.$$

We draw the following consequence.

Lemma 3.1. *Let ψ be the index function for which Assumption 3.1 holds. Then the operators C_0 and T^*T commute. Moreover we have that*

$$\Theta_\psi^2(C_0) = B^*B.$$

Following the last lemma, we set

$$(14) \quad f(s) := \left((\Theta_\psi^2)^{-1}(s) \right)^{1/2}, \quad s > 0.$$

We stress that the function f is an index function, since the function Θ_ψ was one. Moreover, the function Θ_ψ^2 is strictly increasing, such that its inverse is a well defined strictly increasing index function. Finally, as can be drawn from Lemma 3.1, we have that under Assumption 3.1 it holds

$$(15) \quad C_0^{1/2} = f(B^*B).$$

Remark 3.2. We remark the following about Assumption 3.1.

- The case that the operator T is the identity is not covered by this assumption. This would require the function $\psi \equiv 1$, which does not constitute an index function. However, for the subsequent analysis we shall only use Lemma 3.1. As seen from (15) we obtain that $\Theta_\psi(t) = \sqrt{t}$, $t > 0$, in this case.
- If the prior C_0 has eigenvalues with multiplicities higher than one, then by Assumption 3.1 the operator T^*T also needs to have eigenvalues with higher multiplicities, since taking functions of operators preserves or increases the multiplicities of the eigenvalues. This is not realistic, hence one should choose a prior covariance with eigenvalues of multiplicity one. This can be achieved by a slight perturbation of the original choice.

In order to have a handy notation we agree to introduce the following partial ordering between index functions.

Notation. Let f, g be index functions. We say that $f \prec g$ if the quotient g/f is non-decreasing. In other words $f \prec g$ if g decays to zero faster than f .

For bounding the bias below we shall assume that the smoothness of the underlying true data-generating element x^* , is given as a source set with respect to C_0 .

Definition 2 (source set). *There is an index function φ such that*

$$x^* \in A_\varphi := \{x, \quad x = \varphi(C_0)w, \quad \|w\| \leq 1\}.$$

By Lemma 3.1 the source set A_φ can be rewritten as

$$A_\varphi = \{x, \quad x = \varphi(f^2(B^*B))w, \quad \|w\| \leq 1\},$$

with the function f from (14). Furthermore, under Assumption 3.1 the operators C_0 and B^*B commute, and hence the bias representation from (9) simplifies to

$$(16) \quad b_{x^*}(\alpha) = \|r_\alpha(B^*B)s_\alpha(B^*B)x^*\|.$$

Overall, if $x^* \in A_\varphi$ then

$$b_{x^*}(\alpha) \leq \|r_\alpha(B^*B)s_\alpha(B^*B)\varphi(f^2(B^*B))\| = \sup_{0 < t \leq \|B^*B\|} |r_\alpha(t)|s_\alpha(t)\varphi(f^2(t)).$$

We shall bound this in terms of the parameter $\alpha > 0$, which directs us to the notion of a *qualification* of a regularization, see [6], again.

Definition 3 (qualification). *A regularization g_α has qualification φ with constant γ , for an index function φ , if*

$$|r_\alpha(t)|\varphi(t) \leq \gamma\varphi(\alpha), \quad \alpha > 0, \quad 0 < t \leq \|B^*B\|.$$

The following result is a well-known consequence, see [6, Prop. 2.7] again, albeit important for the subsequent analysis. We shall use the partial ordering from Definition 3.1.

Lemma 3.2. *Let g_α be a regularization with index function φ as a qualification (with constant γ). If ψ is an index function for which $\psi \prec \varphi$ then ψ is also a qualification (with constant γ).*

Remark 3.3. As seen from the above analysis of the bias, we shall apply this to the compound function $r_\alpha(t)s_\alpha(t)$, which is related to the compound regularization, obtained by pre-conditioning and Tikhonov regularization. Clearly, it is desirable to bound the bias by a function of α which decays to zero as quickly as possible. It is thus apparent, that a qualification φ of a regularization quantifies its capability to take smoothness, given in terms of source sets, into account.

Example 4 (Tikhonov regularization). Tikhonov regularization has (maximal) qualification $\varphi(t) = t$, $t > 0$. Thus, if for an index function ψ we have that $\psi(t) \prec t$ then ψ is a qualification. In particular, all concave index functions are qualifications of Tikhonov regularization with constant $\gamma = 1$.

Example 5 (spectral cut-off). Spectral cut-off has arbitrary qualification, since $r_\alpha(t) = 0$, $t \geq \alpha$ and $r_\alpha(t) = 1$ elsewhere. Hence

$$r_\alpha(t)\varphi(t) = 0 \leq \varphi(\alpha), \quad t \geq \alpha, \quad \text{and} \quad r_\alpha(t)\varphi(t) \leq \varphi(\alpha), \quad t \leq \alpha.$$

Remark 3.4. We immediately see from (9) that the qualification of the regularization in the bias, can be raised from t (Tikhonov regularization) to t^{k+1} , if the residual function r_α of the regularization used for preconditioning the prior mean has qualification t^k , as is the case for k -fold Tikhonov regularization, see Example 2. If preconditioning is done by spectral cut-off, then the regularization in the bias has arbitrary qualification.

3.2. Bounding the bias. We are now ready to present our bounds for the bias.

Proposition 3.1. *Suppose that $x^* \in A_\varphi$, and that m_α^δ uses a regularization g_α with constant γ_0 bounding the corresponding residual function.*

- (1) *If $\varphi \prec \Theta_\psi^2$, then $b_{x^*}(\alpha) \leq \gamma_0\varphi(f^2(\alpha))$, $\alpha > 0$.*
- (2) *If $\Theta_\psi^2 \prec \varphi$ and if there was no preconditioning, then there are constants $c_1, c_2 > 0$ (depending on x^*, φ, f^2 , and on $\|B^*B\|$) such that $c_1\alpha \leq b_{x^*}(\alpha) \leq c_2\alpha$, $0 < \alpha \leq 1$.*

- (3) If $\Theta_\psi^2 \prec \varphi$ and if $t \mapsto \varphi(f^2(t))/t$ is a qualification for the regularization g_α with constant γ , then $b_{x^*}(\alpha) \leq \gamma \varphi(f^2(\alpha))$, $\alpha > 0$.

Remark 3.5. We mention that the above two cases $\varphi \prec \Theta_\psi^2$ or $\Theta_\psi^2 \prec \varphi$ are nearly disjoint, with $\varphi = \Theta_\psi^2$ being the only common member. Therefore the function Θ_ψ^2 may be viewed as the 'benchmark smoothness'. However, note that the items (1) and (3) do not exhaust all possibilities since the function $\varphi(f^2(t))/t$ may not be a qualification for g_α (in fact it may not even be an index function).

Remark 3.6. We stress that the bounds in item (2) show the saturation phenomenon in the bias if no preconditioning of the prior mean is used: for any sufficiently high smoothness the bias decays with the fixed rate α . In other words, if no preconditioning of the prior is used, the best achievable rate of decay for the bias is linear. Item (3) shows that appropriate preconditioning improves things, since for high smoothness the bias decays at the superlinear rate $\varphi(f^2(\alpha))$.

3.3. The net posterior spread. Here we study the posterior spread, that is, the trace of the posterior covariance from (8), which will be needed for determining the contraction rate. In order to highlight the nature of the spread in the posterior within the assumed Bayesian framework, we make the following definition, for a given equation $z^\delta = Tx + \delta\xi$, with white noise ξ , as considered in (1).

Definition 4 (net posterior spread). *The function*

$$S_{T,C_0}(\alpha) := \text{tr} \left[C_0^{1/2} (\alpha I + B^* B)^{-1} C_0^{1/2} \right], \quad \alpha > 0,$$

is called the net posterior spread.

Notice that with this function we have that $\text{tr} [C^\delta(\alpha)] = \delta^2 S_{T,C_0}(\alpha)$. Moreover, using the cyclic commutativity of the trace, we get that

$$(17) \quad S_{T,C_0}(\alpha) = \text{tr} [(\alpha I + B^* B)^{-1} C_0].$$

With this more convenient representation at hand, we establish some fundamental properties of the net posterior spread, which are crucial for optimizing the convergence rate of SPC in the following subsection.

Lemma 3.3.

- (1) *The function $\alpha \mapsto S_{T,C_0}(\alpha)$ is strictly decreasing and continuous for $\alpha > 0$.*
- (2) $\lim_{\alpha \rightarrow \infty} S_{T,C_0}(\alpha) = 0$, and
- (3) $\lim_{\alpha \rightarrow 0} S_{T,C_0}(\alpha) = \infty$.

3.4. Bounding the squared posterior contraction. It has already been highlighted that the squared posterior contraction as given in (2) is decomposed into the sum of the squared bias, estimation variance and posterior spread, see (3). By Proposition 2.1 we find that

$$b_{x^*}^2(\alpha) + \delta^2 S_{T,C_0}(\alpha) \leq \text{SPC}(\alpha) \leq b_{x^*}^2(\alpha) + ((1 + \gamma_*)^2 + 1) \delta^2 S_{T,C_0}(\alpha).$$

In the asymptotic regime of $\delta \rightarrow 0$, the size of SPC is thus determined by the sum $b_{x^*}^2(\alpha) + \delta^2 S_{T,C_0}(\alpha)$. In § 3.2 we have established bounds for the bias. Here we just constrain to the case where, given that $x^* \in A_\varphi$, the preconditioning is such that the size of the bias is bounded by (a multiple of) $\varphi(f^2(\alpha))$, see Proposition 3.1. Since $b_{x^*}^2(\alpha)$ is bounded by a non-decreasing function of α which decays to zero as $\alpha \searrow 0$, while by Lemma 3.3 the function $S_{T,C_0}(\alpha)$ is strictly decreasing, continuous and onto the positive half-line, the SPC is 'minimized' by the choice of α which balances the bound for the squared bias and the spread. This choice clearly exists and is unique and hence we immediately arrive to our main result.

Theorem 1. *Let φ be any index function, and assume that item (1) or item (3) in Proposition 3.1 hold. Consider the equation*

$$(18) \quad \varphi^2(f^2(\alpha)) = \delta^2 S_{T,C_0}(\alpha).$$

The equation (18) is uniquely solvable, and let $\alpha_ = \alpha_*(\varphi, \delta)$ be the solution. For $x^* \in A_\varphi$ we have that $\text{SPC}(\alpha_*, \delta) = \mathcal{O}(\varphi^2(f^2(\alpha_*)))$ as $\delta \rightarrow 0$.*

The importance of this theorem will become apparent in the next section. In many specific cases, the obtained contraction rates of the SPC correspond to known minimax rates in statistical inverse problems. This can be seen in Propositions 4.2, 4.4, 4.6 and 4.8 below. For general link conditions and general source conditions, minimax rates and in particular lower bounds are scarce. Here we mention the study [14], where the linking function ψ is of power type, and the smoothness function φ is assumed to be concave.

Remark 3.7. As emphasized in Remark 3.6, if no preconditioning is used, the best rate at which the bias can decay is linear. This effect, which is called saturation (of Tikhonov regularization), was discussed in a more general context in regularization theory, and we mention the study [13].

So, if no preconditioning is present, then the left hand side in (18) at best decays as α^2 . We conclude that the best rate of decay of the SPC which can be established without preconditioning is α_*^2 , where α_* is obtained from balancing $\alpha^2 = \delta^2 S_{T,C_0}(\alpha)$. Balancing actually gives (up to some constant) the minimum value, as it was shown in Lemma 2.4 in the same reference.

4. EXAMPLES AND DISCUSSION

We now study several examples, some, which are standard in the literature, and some which exhibit new features. Our aim is to demonstrate the simplicity of our method for deriving rates of posterior contraction and most importantly the benefits of preconditioning the prior.

Before we proceed we stress the following fact, which is not so accurately spelled out in other studies. It is important to distinguish the *degree of ill-posedness of the operator T* which governs equation (1), and which expresses the decay of its singular numbers, from the *degree of ill-posedness of the problem*, which corresponds to the operator T and the solution smoothness, and thus regards the achievable contraction rate. As we will see in § 4.5 below, the problem can have a significantly different degree of ill-posedness than the operator T .

We first consider two examples which concern Sobolev-like smoothness of the truth. We recover the moderately and severely ill-posed problems, as for example studied in [8], and [9, 3], respectively. Then, we consider another two examples which concern analytic-type smoothness of the truth, which to our knowledge have not been studied before. First, we once more study the moderately ill-posed operator problem, which we will see that under analytic-type smoothness of the truth leads to what we call a *mildly ill-posed* problem. Then, we study a problem with severely ill-posed operator, which as we will see, under analytic-type smoothness of the truth leads to a *moderately ill-posed* problem.

In all of the examples, the operators C_0 and T^*T are simultaneously diagonalizable in an orthonormal basis $\{e_j\}$ which is complete in X , C_0 has spectrum that decays as $\{j^{-1-2a}\}$, $a > 0$, while T^*T can either have spectrum that decays polynomially (moderately ill-posed operator case) or exponentially (severely ill-posed operator case).

Notation. Given two positive functions $k, h : \mathbb{R}^+ \rightarrow \mathbb{R}^+$, we use $k \asymp h$ to denote that $k = \mathcal{O}(h)$ and $h = \mathcal{O}(k)$ as $s \rightarrow 0$. Furthermore, the notation $h(s) \gg k(s)$, means that $k(s) = \mathcal{O}(h(s)s^\mu)$ as $s \rightarrow 0$ for some positive power $\mu > 0$.

4.1. Smoothness relative to the prior. In the first two examples, we present posterior contraction rates under the assumption that we have the a priori knowledge that the truth belongs to the Sobolev ellipsoid

$$(19) \quad S^\beta = \{x \in X : \sum_{j=1}^{\infty} j^{2\beta} x_j^2 \leq 1\},$$

for some $\beta > 0$ and where $x_j := \langle x, e_j \rangle$. Relative to C_0 , the index function defining the source set A_φ in Definition 2, is in this case $\varphi(t) = t^{\frac{\beta}{1+2a}}$.

In the third example, we present posterior contraction rates under analytic smoothness of the truth, that is, we assume that we have the a priori knowledge that the truth belongs to the ellipsoid

$$(20) \quad \mathcal{A}^\beta = \{x \in X : \sum_{j=1}^{\infty} e^{2\beta j} x_j^2 \leq 1\},$$

for some $\beta > 0$. In this case, the index function defining the source set A_φ in Definition 2, is $\varphi(t) = \exp(-\beta t^{-\frac{1}{1+2a}})$.

4.2. Moderately ill-posed operator under Sobolev smoothness.

We consider the moderately ill-posed setup studied in [8], in which the operator T^*T has spectrum which decays as $\{j^{-2p}\}$ for some $p \geq 0$, and thus the singular numbers of B^*B decay as $s_j(B^*B) \asymp j^{-(1+2a+2p)}$.

In the present case Assumption 3.1, which expresses the operator T^*T as a function of the prior covariance operator C_0 , is satisfied for $\psi^2(t) = t^{\frac{2p}{1+2a}}$. Next, we find that the function Θ_ψ in (13), which expresses the operator B^*B as a function of C_0 , is given as $\Theta_\psi(t) = t^{\frac{1+2a+2p}{2(1+2a)}}$, hence the benchmark smoothness is $\Theta_\psi^2(t) = t^{\frac{1+2a+2p}{1+2a}}$. Finally, we have that the function f in (14), which expresses C_0 as a function of B^*B is given by $f(s) = s^{\frac{1+2a}{2(1+2a+2p)}}$.

Bounding the bias. We now have all the ingredients required to bound the bias. The following result is an immediate consequence of Proposition 3.1 and the considerations of the previous paragraph.

Proposition 4.1. *Suppose that $x^* \in S^\beta$, for some $\beta > 0$. Then as $\alpha \rightarrow 0$:*

- (1) *If $\beta \leq 1+2a+2p$, and independently of whether preconditioning of the prior is used or not, we have that $b_{x^*}(\alpha) = \mathcal{O}(\alpha^{\frac{\beta}{1+2a+2p}})$;*
- (2) *if $\beta > 1+2a+2p$ and no preconditioning of the prior is used, then $b_{x^*}(\alpha) \asymp \alpha$;*
- (3) *if $\beta > 1+2a+2p$ and m_α^δ uses a regularization g_α with qualification $t^{\frac{\beta-1-2a-2p}{1+2a+2p}}$, then $b_{x^*}(\alpha) = \mathcal{O}(\alpha^{\frac{\beta}{1+2a+2p}})$.*

We stress here that our contribution is item (3). In particular, item (3) implies that if we choose the prior mean m_α^δ using the k -fold Tikhonov regularization filter (cf. Example 2), which has maximal qualification t^k , then for $\beta \leq (k+1)(1+2a+2p)$ we have that $b_{x^*}(\alpha) = \mathcal{O}(\alpha^{\frac{\beta}{1+2a+2p}})$, that is, the saturation in the bias is delayed. If we choose m_α^δ using the spectral cut-off regularization filter, which as we saw in Example 5 has arbitrary qualification, then for any $\beta > 0$, we have that $b_{x^*}(\alpha) = \mathcal{O}(\alpha^{\frac{\beta}{1+2a+2p}})$, that is, there is no saturation in the bias.

Bounding the SPC. To see the impact of this result to the SPC rate, we apply Theorem 1. In order to do so, we first need to calculate the net posterior spread which in this case is such that $S_{T,C_0}(\alpha) \asymp \alpha^{-\frac{1+2p}{1+2a+2p}}$, see [8, Thm 4.1]. Concatenating we get the following result.

Proposition 4.2. *Suppose that $x^* \in S^\beta$, $\beta > 0$. Then as $\delta \rightarrow 0$:*

- (1) if $\beta \leq 1 + 2a + 2p$ and independently of whether preconditioning of the prior is used or not, for $\alpha = \delta^{\frac{2(1+2a+2p)}{1+2p+2\beta}}$ we have that $SPC = \mathcal{O}(\delta^{\frac{4\beta}{1+2\beta+2p}})$;
- (2) if $\beta > 1 + 2a + 2p$ and no preconditioning of the prior is used, then for any choice $\alpha = \alpha(\delta, \beta)$ we have that $SPC \gg \delta^{\frac{4\beta}{1+2\beta+2p}}$;
- (3) if $\beta > 1 + 2a + 2p$ and m_α^δ uses a regularization g_α with qualification $t^{\frac{\beta-1-2a-2p}{1+2a+2p}}$, for $\alpha = \delta^{\frac{2(1+2a+2p)}{1+2p+2\beta}}$ we have that $SPC = \mathcal{O}(\delta^{\frac{4\beta}{1+2\beta+2p}})$.

As before, our contribution is item (3), which in particular implies that if we choose the prior mean m_α^δ using the k -fold Tikhonov regularization filter, then for $\beta \leq (k+1)(1+2a+2p)$ we achieve the optimal (minimax) rate $\delta^{\frac{4\beta}{1+2\beta+2p}}$, that is the saturation in the SPC is also delayed. If we choose m_α^δ using the spectral cut-off regularization filter, then for any $\beta \geq 0$ we achieve the optimal rate $\delta^{\frac{4\beta}{1+2\beta+2p}}$, that is, there is no saturation in the SPC! Note that the optimal scaling of the prior, as a function of the noise level δ , is the same whether we use preconditioning or not. We depict the findings in Figure 1.

4.3. Severely ill-posed operator under Sobolev smoothness.

We now consider the severely ill-posed setup studied in [3, 9], in which the operator T^*T has spectrum which decays as $\{e^{-2qj^b}\}$ for some $q, b > 0$, and thus the singular numbers of B^*B decay as $s_j(B^*B) \asymp j^{-(1+2a)}e^{-2qj^b}$.

In this case Assumption 3.1, which expresses the operator T^*T as a function of the prior covariance operator C_0 , is satisfied for $\psi^2(t) = \exp(-2qt^{-\frac{b}{1+2a}})$. Next, we find that the function Θ_ψ in (13), which expresses the operator B^*B as a function of C_0 , is given as $\Theta_\psi(t) = t^{\frac{1}{2}} \exp(-qt^{-\frac{b}{1+2a}})$, and hence the benchmark smoothness is $\Theta_\psi^2(t) = t \exp(-2qt^{-\frac{b}{1+2a}})$. Finally, we have that as $t \rightarrow 0$, the function f in (14) which expresses C_0 as a function of B^*B behaves as $f(s) \sim (\log(s^{-\frac{1}{2q}}))^{-\frac{1+2a}{2b}}$, see Lemma 5.1 in Section 5.

Bounding the bias. In this example we have that $\Theta_\psi^2(t)$ decays exponentially, while $\varphi(t)$ polynomially, hence for any Sobolev-like smoothness of the truth β , it holds $\varphi \prec \Theta_\psi^2$. In other words, even without preconditioning there is no saturation in the bias and we are always in case (1) in Proposition 3.1. However, our theory still works and we can easily derive the rate for the bias and SPC. The next result follows immediately from the considerations in the previous paragraph and Proposition 3.1.

Proposition 4.3. *Suppose that $x^* \in S^\beta$, $\beta > 0$. Then independently of whether preconditioning of the prior is used or not, we have that $b_{x^*}(\alpha) = \mathcal{O}((\log(\alpha^{-1}))^{-\frac{\beta}{b}})$, as $\alpha \rightarrow 0$.*

Bounding the SPC. We now apply Theorem 1 in order to calculate the SPC rate. Again, we first need to calculate the net posterior spread, which in this case is such that $S_{T,C_0}(\alpha) \asymp \frac{1}{\alpha}(\log(\alpha^{-1}))^{-\frac{2a}{b}}$, see [3, Thm 4.2]. We prove the following result, which agrees with [9, Thm 2.1] and [3, Thm 4.3].

Proposition 4.4. *Suppose that $x^* \in S^\beta$, $\beta > 0$. Then independently of whether preconditioning of the prior is used or not, for any $\sigma > 0$, any parameter choice rule $\alpha = \alpha(\delta)$ such that $\delta^2(\log(\delta^{-2}))^{\frac{2\beta-2a}{b}} \leq \alpha \leq \delta^{2\sigma}$, gives the rate $SPC = \mathcal{O}((\log(\delta^{-2}))^{-\frac{2\beta}{b}})$, as $\delta \rightarrow 0$.*

4.4. Moderately ill-posed operator under analytic smoothness.

We now consider the moderately ill-posed operator setup studied in § 4.2 with the difference that here we assume that we have the a priori knowledge that the truth has a certain analytic smoothness. The functions ψ , Θ_ψ and f which have to do with the relationship between the forward operator and the prior covariance are as in § 4.2, but the function φ which describes analytic smoothness of the truth as in (20), is now $\varphi(t) = \exp(-\beta t^{-\frac{1}{1+2a}})$. In particular, since φ is exponential while the benchmark smoothness Θ_ψ^2 is of power type, we are always in the high smoothness case $\Theta_\psi^2 \prec \varphi$.

Bounding the bias. The following is an immediate consequence of Proposition 3.1 and the considerations in the previous paragraph.

Proposition 4.5. *Suppose that $x^* \in \mathcal{A}^\beta$, for some $\beta > 0$. Then as $\alpha \rightarrow 0$:*

- (1) *if no preconditioning is used, $b_{x^*}(\alpha) \asymp \alpha$;*
- (2) *if m_α^δ uses a regularization g_α with qualification $\exp(-\beta t^{-1})$, then we have that $b_{x^*}(\alpha) = \mathcal{O}(\exp(-\beta \alpha^{-\frac{1}{1+2a+2p}}))$.*

Remark 4.1. If no preconditioning is used, the bias convergence rate is always saturated. The qualification as formulated in item (2) is a *sufficient condition*, while the actual form can be calculated easily. The given form highlights that exponential type qualification is required to overcome the limitation of the power type prior covariance in order to treat analytic smoothness. We stress here that such qualification is hard to achieve. For example, iterated Tikhonov can never achieve such exponential qualification, while even Landweber iteration which has qualification t^ν , for any $\nu > 0$, only achieves this qualification for values β which are not too big. On the other hand, $\exp(-\beta t^{-1})$ is a qualification for spectral cut-off for any positive value of β .

Bounding the SPC. We again apply Theorem 1 in order to calculate the SPC rate. The net posterior spread is as in § 4.2, $S_{T,C_0}(\alpha) \asymp \alpha^{-\frac{1+2p}{1+2a+2p}}$. We prove the following result, using the convention from Definition 4.

Proposition 4.6. *Suppose that $x^* \in \mathcal{A}^\beta$, $\beta > 0$. Then as $\delta \rightarrow 0$:*

- (1) *if no preconditioning of the prior is used, then for any choice $\alpha = \alpha(\delta, \beta)$ we have that $SPC \gg \delta^2$;*
- (2) *if m_α^δ uses a regularization g_α with qualification $\exp(-\beta t^{-1})$, for $\alpha = (\log(\delta^{-1/\beta}))^{-(1+2a+2p)}$ we have that $SPC = \mathcal{O}(\delta^2(\log(\delta^{-1}))^{1+2p})$.*

Remark 4.2. We stress that according to item (1), without preconditioning we have that δ^2/SPC decays at an algebraic rate, while the optimal achievable (also minimax) rate is of power two up to some logarithmic factor. Since the optimal achievable rate in this case is of power two up to logarithmic factors, it is reasonable to call such problems *mildly ill-posed*, as they are almost well-posed.

4.5. Severely ill-posed operator under analytic smoothness.

We now consider the severely ill-posed operator setup studied in § 4.3 with the difference that here we assume that we have the a priori knowledge that the truth has a certain analytic smoothness. For simplicity, we concentrate on the case $b = 1$, which corresponds for example to the Cauchy problem for the Helmholtz equation, see [3, Section 5] for details.

The functions ψ, Θ_ψ and f which have to do with the relationship between the forward operator and the prior covariance are as in § 4.3 for the value $b = 1$, but the function φ which describes analytic smoothness of the truth as in (20), is now $\varphi(t) = \exp(-\beta t^{-\frac{1}{1+2a}})$. In particular, since both φ and the benchmark smoothness Θ_ψ^2 are exponential, unlike § 4.3 we now have a saturation phenomenon.

Bounding the bias. The following is an immediate consequence of Proposition 3.1 and the considerations in the previous paragraph.

Proposition 4.7. *Suppose that $x^* \in \mathcal{A}^\beta$, for some $\beta > 0$. Then as $\alpha \rightarrow 0$:*

- (1) *if $\beta \leq 2q$ and independently of whether preconditioning of the prior is used or not, we have that $b_{x^*}(\alpha) = \mathcal{O}(\alpha^{\frac{\beta}{2q}})$;*
- (2) *if $\beta > 2q$ and no preconditioning is used $b_{x^*}(\alpha) \asymp \alpha$;*
- (3) *if $\beta > 2q$ and m_α^δ uses a regularization g_α with qualification $t^{\frac{\beta-2q}{2q}}$, then we have that $b_{x^*}(\alpha) = \mathcal{O}(\alpha^{\frac{\beta}{2q}})$.*

The benefits of preconditioning are once more clear and can be seen in item (3). If for example we choose the prior mean m_α^δ using the k -fold Tikhonov regularization filter, then for $\beta \leq (k+1)2q$ we have that $b_{x^*}(\alpha) = \mathcal{O}(\alpha^{\frac{\beta}{2q}})$, that is the saturation in the bias is delayed. If we use spectral cut-off, then there is no saturation at all.

Bounding the SPC. We again apply Theorem 1 in order to calculate the SPC rate. The net posterior spread is as in § 4.3, $S_{T,C_0}(\alpha) \asymp \frac{1}{\alpha}(\log(\alpha^{-1}))^{-2a}$. We prove the following result.

Proposition 4.8. *Suppose that $x^* \in \mathcal{A}^\beta$, $\beta > 0$. Then as $\delta \rightarrow 0$:*

- (1) *If $\beta \leq 2q$ and independently of whether preconditioning of the prior is used or not, for $\alpha = \delta^{\frac{2q}{\beta+q}}$ we have that $SPC = \mathcal{O}(\delta^{\frac{2\beta}{\beta+q}})$;*
- (2) *if $\beta > 2q$ and no preconditioning of the prior is used, then for any choice $\alpha = \alpha(\delta, \beta)$ we have that $SPC \gg \delta^{\frac{2\beta}{\beta+q}}$;*
- (3) *if $\beta > 2q$ and m_α^δ uses a regularization g_α with qualification $t^{\frac{\beta-2q}{2q}}$, for $\alpha = \delta^{\frac{2q}{\beta+q}}$ we have that $SPC = \mathcal{O}(\delta^{\frac{2\beta}{\beta+q}})$.*

The benefits of preconditioning can again be seen in item (3). If for example we choose the prior mean m_α^δ using the k -fold Tikhonov regularization filter, then for $\beta \leq (k+1)2q$ we achieve the optimal (minimax) rate $\delta^{\frac{2\beta}{\beta+q}}$, that is the saturation in the SPC is delayed. If we use spectral cut-off, then there is no saturation at all. Note again that the optimal scaling of the prior, as a function of the noise level δ , is the same whether we use preconditioning or not.

As anticipated, the features of this example, and in particular the polynomial rates of convergence, are characteristic of moderately ill-posed problems.

4.6. Summary and discussion. We succinctly summarize the above examples, in which we confined to power-type decay of the spectrum of the prior C_0 , that is, $s_j(C_0) \asymp j^{-(1+2a)}$, $j = 1, 2, \dots$, for some $a > 0$.

First in § 4.2 and § 4.3, we specified the solution element to belong to some Sobolev-type ball as in (19), characterized by $\beta > 0$. The distinction between moderately and severely ill-posed problems then comes from the decay of the singular numbers of the operator T governing equation (1). We outline the previous results in Table 1.

Then in § 4.4 and § 4.5, we considered analytic type smoothness of the truth as in (20), again characterized by $\beta > 0$. As commented earlier on, to our knowledge we are the first to study these examples. Our findings show that the overall problem degree of ill-posedness can be significantly different than the degree of ill-posedness of the operator. We outline the results in Table 2.

The rates exhibited in Tables 1 and 2, correspond to the minimax rates as given in [4, Tbl. 1].

		$s_j(T^*T) \asymp j^{-2p}$	$s_j(T^*T) \asymp e^{-2qj^b}$
link	ψ	$t^{p/(1+2a)}$	$\exp(-2qt^{-b/(1+2a)})$
benchmark	Θ_ψ^2	$t^{(1+2a+2p)/(1+2a)}$	$t \exp(-2qt^{-b/(1+2a)})$
saturation	$\varphi = \Theta_\psi^2$	$\beta = 1 + 2a + 2p$	always $\varphi \prec \Theta_\psi^2$
contraction	SPC	$\delta^{4\beta/(1+2a+2p)}$	$\log^{-2\beta/b}(1/\delta)$

TABLE 1. Outline of SPC rates for Sobolev-type smoothness of the truth, $\varphi(t) = t^{\beta/(1+2a)}$, $t > 0$.

		$s_j(T^*T) \asymp j^{-2p}$	$s_j(T^*T) \asymp e^{-2qj^b}$
link	ψ	$t^{p/(1+2a)}$	$\exp(-2qt^{-b/(1+2a)})$
benchmark	Θ_ψ^2	$t^{(1+2a+2p)/(1+2a)}$	$t \exp(-2qt^{-b/(1+2a)})$
saturation	$\varphi = \Theta_\psi^2$	always $\Theta_\psi^2 \prec \varphi$	$\beta = 2q$
contraction	SPC	$\delta^2 \log^{1+2p}(1/\delta)$	$\delta^{2\beta/(\beta+q)}$

TABLE 2. Outline of SPC rates for analytic-type smoothness of the truth, $\varphi(t) = \exp(-\beta t^{-\frac{1}{1+2a}})$, $t > 0$.

5. PROOFS AND AUXILIARY RESULTS

Proof of Lemma 2.1. We first express the element x_α^δ in terms of z^δ .

$$\begin{aligned}
x_\alpha^\delta &= C_0^{1/2} (\alpha I + B^* B)^{-1} B^* z^\delta + C_0^{1/2} s_\alpha(B^* B) C_0^{-1/2} m_\alpha^\delta \\
&= C_0^{1/2} (\alpha I + B^* B)^{-1} B^* z^\delta + C_0^{1/2} s_\alpha(B^* B) g_\alpha(B^* B) B^* z^\delta \\
&= C_0^{1/2} [(\alpha I + B^* B)^{-1} + s_\alpha(B^* B) g_\alpha(B^* B)] B^* z^\delta.
\end{aligned}$$

We notice that

$$(\alpha I + B^* B)^{-1} + s_\alpha(B^* B) g_\alpha(B^* B) = (\alpha I + B^* B)^{-1} (I + \alpha g_\alpha(B^* B)).$$

The expectation of the posterior mean with respect to the distribution generating z^δ when x^* is given, is thus

$$\mathbb{E}^{x^*} x_\alpha^\delta = C_0^{1/2} [(\alpha I + B^* B)^{-1} (I + \alpha g_\alpha(B^* B))] B^* B C_0^{-1/2} x^*.$$

For the next calculations we shall use that

$$\begin{aligned} I - (\alpha I + B^* B)^{-1} (I + \alpha g_\alpha(B^* B)) B^* B \\ &= (\alpha I + B^* B)^{-1} \alpha (I - g_\alpha(B^* B) B^* B) \\ &= s_\alpha(B^* B) r_\alpha(B^* B). \end{aligned}$$

Therefore we rewrite

$$\begin{aligned} x^* - \mathbb{E}^{x^*} x_\alpha^\delta &= C_0^{1/2} [I - (\alpha I + B^* B)^{-1} (I + \alpha g_\alpha(B^* B)) B^* B] C_0^{-1/2} x^* \\ &= C_0^{1/2} s_\alpha(B^* B) r_\alpha(B^* B) C_0^{-1/2} x^*, \end{aligned}$$

which proves the first assertion. The variance is $\mathbb{E}^{x^*} \|x_\alpha^\delta - \mathbb{E}^{x^*} x_\alpha^\delta\|^2$, and this can be written as in (10), by using similar reasoning as for the bias term. \square

Proof of Proposition 2.1. We notice that $\|I + \alpha g_\alpha(B^* B)\| \leq 1 + \gamma_*$, which gives

$$\begin{aligned} V^\delta(\alpha) &= \delta^2 \operatorname{tr} [(I + \alpha g_\alpha(B^* B))^2 (\alpha I + B^* B)^{-2} B^* B C_0] \\ &\leq \delta^2 (1 + \gamma_*)^2 \operatorname{tr} [(\alpha I + B^* B)^{-2} B^* B C_0] \end{aligned}$$

Since $\|(\alpha I + B^* B)^{-1} B^* B\| \leq 1$ we see that

$$V^\delta(\alpha) \leq (1 + \gamma_*)^2 \delta^2 \operatorname{tr} [(\alpha I + B^* B)^{-1} C_0] = (1 + \gamma_*)^2 \operatorname{tr} [C^\delta(\alpha)],$$

and the proof is complete. \square

Proof of Lemma 3.1. Since C_0 has finite trace, it is compact, and we use the eigenbasis (arranged by decreasing eigenvalues) u_j , $j = 1, 2, \dots$. Under Assumption 3.1 this is also the eigenbasis for $T^* T$. If t_j , $j = 1, 2, \dots$ denote the eigenvalues then we see that

$$T^* T = \sum_{j=1}^{\infty} \tau_j u_j \otimes u_j.$$

Correspondingly, $C_0 = \sum_{j=1}^{\infty} (\psi^2)^{-1} (\tau_j) u_j \otimes u_j$, which gives the first assertion. Moreover, the latter representation yields that

$$C_0^{1/2} = \sum_{j=1}^{\infty} \left((\psi^2)^{-1} (\tau_j) \right)^{1/2} u_j \otimes u_j,$$

such that

$$\begin{aligned}
B^*B &= C_0^{1/2} T^* T C_0^{1/2} \\
&= \sum_{j=1}^{\infty} \left((\psi^2)^{-1}(\tau_j) \right)^{1/2} \tau_j \left((\psi^2)^{-1}(\tau_j) \right)^{1/2} u_j \otimes u_j \\
&= \sum_{j=1}^{\infty} \left((\psi^2)^{-1}(\tau_j) \right) \tau_j u_j \otimes u_j \\
&= \sum_{j=1}^{\infty} \psi^2 \left(\left((\psi^2)^{-1}(\tau_j) \right) \right) \left((\psi^2)^{-1}(\tau_j) \right) u_j \otimes u_j \\
&= \sum_{j=1}^{\infty} \Theta_{\psi}^2 \left((\psi^2)^{-1}(\tau_j) \right) u_j \otimes u_j \\
&= \Theta_{\psi}^2(C_0),
\end{aligned}$$

and the proof is complete. \square

Proof of Proposition 3.1. For the first item (1), we notice that $\varphi \prec \Theta_{\psi}^2$ if and only if $\varphi(f^2(t)) \prec t$. The linear function $t \mapsto t$ is a qualification of Tikhonov regularization with constant $\gamma = 1$. Thus, by Lemma 3.2 we have

$$b_{x^*}(\alpha) \leq \|r_{\alpha}(B^*B)\| \|s_{\alpha}(B^*B)\varphi(f^2(B^*B))\| \leq \gamma_0 \varphi(f^2(\alpha)),$$

which completes the proof for this case. For item (2), we have that

$$b_{x^*}(\alpha) = \|s_{\alpha}(B^*B)x^*\|.$$

For any $0 < \alpha \leq 1$, we have $\alpha + t \leq 1 + t$, hence

$$b_{x^*}(\alpha) = \alpha \|(\alpha I + B^*B)^{-1}x^*\| \geq \alpha \|(I + B^*B)^{-1}x^*\|.$$

We conclude that there exists a constant $c_1 = c_1(x^*, \|B^*B\|)$, such that for small α it holds

$$b_{x^*}(\alpha) \geq c_1 \alpha.$$

On the other hand, since $t \prec \varphi(f^2(t))$, there exists a constant $c_2 > 0$ which depends only on the index functions φ , f and on $\|B^*B\|$, such that

$$b_{x^*}(\alpha) = \alpha \|(\alpha I + B^*B)^{-1}x^*\| \leq \alpha \|(B^*B)^{-1}\varphi(f^2(B^*B))w\| \leq c_2 \alpha.$$

For item (3), we have that

$$\begin{aligned}
b_{x^*}(\alpha) &\leq \|r_{\alpha}(B^*B)s_{\alpha}(B^*B)\varphi(f^2(B^*B))\| \\
&\leq \|s_{\alpha}(B^*B)B^*B\| \|r_{\alpha}(B^*B)\varphi(f^2(B^*B))(B^*B)^{-1}\| \\
&\leq \alpha \gamma \frac{\varphi(f^2(\alpha))}{\alpha} = \gamma \varphi(f^2(\alpha)),
\end{aligned}$$

and the proof is complete. \square

Proof of Lemma 3.3. The continuity is clear. For the monotonicity we use the representation (17) to get

$$\begin{aligned} S_{T,C_0}(\alpha) - S_{T,C_0}(\alpha') &= \operatorname{tr} [(\alpha I + B^* B)^{-1} C_0] - \operatorname{tr} [(\alpha' + B^* B)^{-1} C_0] \\ &= \operatorname{tr} [(\alpha I + B^* B)^{-1} (\alpha' - \alpha) (\alpha' + B^* B)^{-1} C_0] \\ &= (\alpha' - \alpha) \operatorname{tr} [(\alpha I + B^* B)^{-1} (\alpha' + B^* B)^{-1} C_0]. \end{aligned}$$

The trace on the right hand side is positive. Indeed, if (s_j^2, u_j, u_j) denotes the singular value decomposition of $B^* B$ then this trace can be written as

$$\operatorname{tr} [(\alpha I + B^* B)^{-1} (\alpha' + B^* B)^{-1} C_0] = \sum_{j=1}^{\infty} \frac{1}{\alpha + s_j^2} \frac{1}{\alpha' + s_j^2} \langle C_0 u_j, u_j \rangle,$$

where the right hand side is positive since the operator C_0 is positive definite. Thus, if $\alpha < \alpha'$ then $S_{T,C_0}(\alpha) - S_{T,C_0}(\alpha')$ is positive, which proves the first assertion.

The proof of the second assertion is simple, and hence omitted. To prove the last assertion we use the partial ordering of self-adjoint operators in Hilbert space, that is, we write $A \leq B$ if $\langle Ax, x \rangle \leq \langle Bx, x \rangle$, $x \in X$, for two self-adjoint operators A and B . Plainly, with $a := \|T^* T\|$, we have that $T^* T \leq aI$. Multiplying from the left and right by $C_0^{1/2}$ this yields $B^* B \leq aC_0$, and thus for any $\alpha > 0$ that $\alpha I + B^* B \leq \alpha I + aC_0$. The function $t \mapsto -1/t$, $t > 0$ is operator monotone, which gives $(\alpha I + aC_0)^{-1} \leq (\alpha I + B^* B)^{-1}$. Multiplying from the left and right by $C_0^{1/2}$ again, we arrive at

$$C_0^{1/2} (\alpha I + aC_0)^{-1} C_0^{1/2} \leq C_0^{1/2} (\alpha I + B^* B)^{-1} C_0^{1/2}.$$

This in turn extends to the traces and gives that

$$\operatorname{tr} [C_0^{1/2} (\alpha I + aC_0)^{-1} C_0^{1/2}] \leq \operatorname{tr} [C_0^{1/2} (\alpha I + B^* B)^{-1} C_0^{1/2}] = S_{T,C_0}(\alpha).$$

Now, let us denote by t_j , $j \in \mathbb{N}$, the singular numbers of C_0 , then we can bound

$$S_{T,C_0}(\alpha) \geq \operatorname{tr} [(\alpha I + aC_0)^{-1} C_0] \geq \sum_{t_j \geq \alpha/a} \frac{t_j}{\alpha + at_j} \geq \frac{1}{2a} \# \left\{ j, t_j \geq \frac{\alpha}{a} \right\}.$$

If $S_{T,C_0}(\alpha)$ were uniformly bounded from above, then there would exist a finite natural number, say N , such that $t_N \geq \frac{\alpha}{a} > t_{N+1}$, for $\alpha > 0$ small enough. But this would imply that $t_{N+1} = 0$, which contradicts the assumption that C_0 is positive definite. \square

Lemma 5.1. *For $t > 0$ let $\Theta_{\Psi}^2(t) = t \exp(-2qt^{-\frac{b}{1+2a}})$, for some $q, b, a > 0$. Then for small s we have $(\Theta_{\Psi}^2)^{-1}(s) \sim (\log s^{-\frac{1}{2q}})^{-\frac{1+2a}{b}}$.*

Proof. Let

$$(21) \quad s = \Theta_{\Psi}^2(t) > 0$$

and observe that t is small if and only if s is small. Applying [3, Lemma 4.5] for $x = t^{-1}$ we get the result. \square

Proof of Proposition 4.4. In this example the explicit solution of Eq. (18) in Theorem 1 is more difficult. However, as discussed in § 3.4, it suffices to asymptotically balance the squared bias and the posterior spread using an appropriate parameter choice $\alpha = \alpha(\delta)$. Indeed, under the stated choice of α the squared bias is of order

$$(\log(\alpha^{-1}))^{-\frac{2\beta}{b}} \leq \sigma^{-\frac{2\beta}{b}} \log(\delta^{-2})^{-\frac{2\beta}{b}}$$

while the posterior spread term is of order

$$\frac{\delta^2}{\alpha} (\log(\alpha^{-1}))^{-\frac{2a}{b}} \leq \log(\delta^{-2})^{-\frac{2\beta}{b}}.$$

\square

Proof of Proposition 4.6. According to the considerations in Remark 3.7, it is straightforward to check that without preconditioning the best SPC rate that can be established is $\delta^{\frac{4+8a+8p}{3+4a+6p}}$ which proves item (1). In the preconditioned case, the explicit solution of Eq. (18) in Theorem 1, which in this case has the form

$$\exp(-2\beta\alpha^{-\frac{1}{1+2a+2p}}) = \delta^2\alpha^{-\frac{1+2p}{1+2a+2p}},$$

is again difficult. However, as discussed in § 3.4, it suffices to asymptotically balance the squared bias and the posterior spread using an appropriate parameter choice $\alpha = \alpha(\delta)$. Indeed, using [3, Lem 4.5] we have that the solution to the above equation behaves asymptotically as the stated choice of α , and substitution gives the claimed rate. \square

Proof of Proposition 4.8. We begin with items (1) and (3). The explicit solution of Eq. (18) in Theorem 1, which in this case has the form

$$\alpha^{\frac{\beta}{q}} = \frac{\delta^2}{\alpha} (\log(\alpha^{-1}))^{-2a},$$

is difficult. As discussed in § 3.4, it suffices to asymptotically balance the squared bias and the posterior spread using an appropriate parameter choice $\alpha = \alpha(\delta)$. Indeed, under the stated choice of α both quantities are bounded from above by $\delta^{\frac{2\beta}{\beta+q}}$. For item (2), according to the considerations in Remark 3.7, it is straightforward to check that without preconditioning the best SPC rate that can be established is $\delta^{\frac{4q}{\beta+q}}$. \square

REFERENCES

1. S. Agapiou, J. M. Bardsley, O. Papaspiliopoulos, and Andrew M. Stuart, *Analysis of the Gibbs sampler for hierarchical inverse problems*, arXiv preprint arXiv:1311.1138 (2013).
2. S. Agapiou, S. Larsson, and A. M. Stuart, *Posterior contraction rates for the Bayesian approach to linear ill-posed inverse problems*, Stochastic Process. Appl. **123** (2013), no. 10, 3828–3860. MR 3084161
3. S. Agapiou, A. M. Stuart, and Y. X. Zhang, *Bayesian posterior contraction rates for linear severely ill-posed inverse problems*, J. Inverse Ill-Posed Probl. **22** (2014), no. 3, 297–321. MR 3215928
4. L. Cavalier, *Nonparametric statistical inverse problems*, Inverse Problems **24** (2008), no. 3, 034004, 19. MR 2421941 (2009j:62109)
5. M. Dashti and A. M. Stuart, *The Bayesian Approach To Inverse Problems*, ArXiv e-prints (2013).
6. B. Hofmann and P. Mathé, *Analysis of profile functions for general linear regularization methods*, SIAM J. Numer. Anal. **45** (2007), no. 3, 1122–1141 (electronic). MR 2318806 (2008i:65117)
7. B. T. Knapik, B. T. Szabó, A. W. van der Vaart, and J. H. van Zanten, *Bayes procedures for adaptive inference in inverse problems for the white noise model*, arXiv:1209.3628v2 [math.ST] (2013).
8. B. T. Knapik, A. W. van der Vaart, and J. H. van Zanten, *Bayesian inverse problems with Gaussian priors*, Ann. Statist. **39** (2011), no. 5, 2626–2657. MR 2906881
9. ———, *Bayesian recovery of the initial condition for the heat equation*, Comm. Statist. Theory Methods **42** (2013), no. 7, 1294–1313. MR 3031282
10. M. S. Lehtinen, L. Päivärinta, and E. Somersalo, *Linear inverse problems for generalised random variables*, Inverse Problems **5** (1989), no. 4, 599–612. MR 1009041 (91g:65129)
11. K. Lin, S. Lu, and P. Mathé, *Oracle-type posterior contraction rates in Bayesian inverse problems*, submitted, 2014.
12. A. Mandelbaum, *Linear estimators and measurable linear transformations on a Hilbert space*, Z. Wahrsch. Verw. Gebiete **65** (1984), no. 3, 385–397. MR 731228 (85h:62030)
13. P. Mathé, *Saturation of regularization methods for linear ill-posed problems in Hilbert spaces*, SIAM J. Numer. Anal. **42** (2004), no. 3, 968–973 (electronic). MR 2112789 (2005m:65093)
14. P. Mathé and S. V. Pereverzev, *Regularization of some linear ill-posed problems with discretized random noisy data*, Math. Comp. **75** (2006), no. 256, 1913–1929 (electronic). MR 2240642 (2008b:65072)
15. B. T. Szabó, A. W. van der Vaart, and J. H. van Zanten, *Empirical Bayes scaling of Gaussian priors in the white noise model*, Electron. J. Stat. **7** (2013), 991–1018. MR 3044507

E-mail address: `Sergios.Agapiou@warwick.ac.uk`

MATHEMATICS INSTITUTE, UNIVERSITY OF WARWICK, COVENTRY CV4 7AL,
UNITED KINGDOM

E-mail address: `peter.mathe@wias-berlin.de`

WEIERSTRASS INSTITUTE FOR APPLIED ANALYSIS AND STOCHASTICS, MOHREN-
STRASSE 39, 10117 BERLIN, GERMANY